

SYSTEM AND METHOD FOR COMMUNICATING ON A VIRTUAL RING IN AN INTERNET PROTOCOL NETWORK

Field of the invention

The present invention relates to communication on digital networks, and more particularly to a system and a method for creating a virtual ring between nodes in a an Internet Protocol (IP) network and for multicasting datagrams to nodes part of this virtual ring.

Background of the invention

In the present description, the term "Network" designates an ordinary network, based on the Internet Protocol (IP) technology. This network can be a Local Area Network (LAN), but also an Enterprise (private) Intranet or even the (public) Internet. The term "Node" designates the computer systems in the network routing the communications, such as routers, and, also, the computer systems exchanging information on the network, such as workstations and servers.

In a network, nodes must be able to exchange information with other nodes of a same group. For instance, the broadcast of a same information to multiple nodes located in different locations is called "Multicast". In a group of N nodes called a Multicast group as illustrated in Figure 1, each node (101) needs to communicate with the $(N-1)$ other nodes. To do this, each node establishes a session with each other node (100). Usually in IP networks, the Transmission Control Protocol (TCP) is used to communicate between nodes because this protocol allows a reliable transport of data through sessions and takes care of the flow control. This is not the case with the transport protocol called UDP (User Datagram Protocol) which is based on the best effort and which does not provide any session mechanism.

If a node, within a group of N nodes, wants to communicate information to all the other nodes of its group, it requires $N-1$ TCP sessions. If all the nodes need to

communicate together in a full mesh configuration, $Nx(N-1)/2$ TCP sessions are required. It is important to note that since a TCP session is bidirectional, the required number of sessions is $Nx(N-1)/2$ and not $Nx(N-1)$.

The number of sessions can be considerable in a network comprising hundreds or 5 thousands of nodes. It can result in an important overhead with a significant impact in term of bandwidth consumption in the network and resource (data processing and memory) utilisation in each node. In each node, the establishment of the TCP sessions requires data processing resources and the maintenance of these TCP sessions requires memory in particular to store the context of the TCP sessions (TCP Control 10 Block).

In absence of synchronisation at the application level, the nodes can exchange the same piece of information on all the TCP sessions at the same time (communication any to any). This is bandwidth consuming at the network level and resource consuming at the level of each node. An example of this scenario is the exchange of 15 routing information between routers. Each router broadcast routing information to the other routers either periodically or when a change occurs, depending on the routing protocol used in the network. Another example is the synchronisation of multiple servers in a distributed database.

Several solutions exist to limit the number of sessions between nodes. A solution 20 illustrated in Figure 2, is to select a "Rendezvous Point", or a central node, to which all other nodes are connected. The central node (200) is responsible for distributing the information to all the other nodes in the network. This configuration called "Star network" reduces the number of connections ($N-1$ sessions) but the main drawback is due to the fact that the central node is the weakest point of the network. Generally, 25 the central node is duplicated by means of a backup central node (201). This configuration called "Dual star network", requires $(N - 1) + (N - 2)$ connections.

Note: the central node (200) is connected to all other nodes including the backup central node (201). The result is the establishment of $N-1$ TCP sessions. The addition of a second star configuration based on the backup central node (201) 30 requires another $N-1$ TCP sessions. However, since a TCP session already exists

between central node (200) and backup central node (201), this session does not need to be duplicated. In conclusion, the number of sessions required in a dual star configuration is $(N-1) + (N-2) = 2xN-3$

Objects of the invention

5 It is an object of the invention to reduce the bandwidth utilisation in an IP network comprising inter-communicating nodes.

It is another object of the invention to reduce the resource consumption of inter-communicating nodes.

It is a further object of the invention to define several groups of inter-communicating
10 nodes in an IP network.

Summary of the invention

The present invention is directed to a method, system and computer program as defined in independent claims, to use in a node within a network comprising a transport layer protocol providing end to end data transfer, for multicasting
15 datagrams on a virtual ring, each node on the virtual ring being logically connected according to the network transport layer protocol to two and only two neighbor nodes through virtual connections, an upstream neighbor node and a downstream neighbor node.

The method comprises the steps of:

20 • sending a virtual ring datagram to the downstream neighbor node on the virtual ring; said virtual ring datagram comprising:

- a virtual ring identifier;
- means for identifying the node originator of the virtual ring datagram;
- data;

25 when a datagram is received;

- identifying the received datagram;

if the received datagram is a token:

- identifying the virtual ring;
- checking whether the token is valid or not;
- if the token is valid, forwarding the token to the downstream neighbor node on the identified virtual ring.

5 if the received datagram is a virtual ring datagram:

- identifying the virtual ring;
- checking the node originator of the received virtual ring datagram;

if the received virtual ring datagram has not been locally originated:

- processing data comprised in said virtual ring datagram;

10 • forwarding said virtual ring datagram to the downstream neighbor node on the identified virtual ring;

if the received virtual ring datagram has been locally originated:

- removing the virtual ring datagram from the virtual ring.

Further embodiments of the invention are provided in the appended dependent

15 claims.

The foregoing, together with other objects, features, and advantages of this invention can be better appreciated with reference to the following specification, claims and drawings.

Brief description of the drawings

20 The new and inventive features believed characteristics of the invention are set forth in the appended claims. The invention itself, however, as well as a preferred mode of use, further objects and advantages thereof, will best be understood by reference to the following detailed description of an illustrative detailed embodiment when read in conjunction with the accompanying drawings, wherein:

25 • **Figure 1** shows an example of "Full mesh network".

• **Figure 2** shows an example of "Star network".

- **Figure 3** shows an example of "Virtual Ring network" according to the present invention.
- **Figure 4** shows how a token is forwarded from node to node on a Virtual Ring according to the present invention.

5 • **Figure 5** shows the token message according to the present invention.

- **Figure 6** shows how a new node is inserted into the Virtual Ring according to the present invention.
- **Figure 7** shows the result of a new node insertion according to the present invention.

10 • **Figure 8** shows the solicited removal of a node according to the present invention.

- **Figure 9** shows the loss of a node according to the present invention.
- **Figure 10** shows the result of a reconfiguration after the loss of a node according to the present invention.

15 • **Figure 11** describes the algorithm executed by a node when this node receives the token according to the present invention.

- **Figure 12** describes the algorithm executed in the Virtual Ring Manager at receipt of the token according to the present invention.

20 • **Figure 13** describes the algorithm executed in a node in view of inserting this node into the Virtual Ring according to the present invention.

- **Figure 14** describes the algorithm executed in a node in view of removing this node from the Virtual Ring according to the present invention.

- **Figure 15** illustrates the algorithm executed in a node when a neighbor node has been inserted or removed according to the present invention.
- **Figure 16** illustrates the algorithm executed in the Virtual Ring Manager when a node is inserted or removed from the Virtual Ring according to the present invention.
- **Figure 17** illustrates the Node Insertion process according to the present invention.
- **Figure 18** illustrates the Solicited Node Removal process according to the present invention.
- 10 • **Figure 19** illustrates the unsolicited Node Removal process according to the present invention.

Preferred embodiment of the invention

The present invention discloses a network topology based on a virtual ring. The N nodes of the network that need to communicate together, are logically/virtually 15 connected according to a virtual ring, each node communicating with two and only two neighbor nodes: an upstream neighbor node and a downstream neighbor node.

Although the present invention applies to any types of nodes, this invention is particularly interesting when several nodes need to exchange a same piece of information between them.

Description of the invention

Several virtual rings can be implemented on a same physical network, each virtual ring allowing a subset of nodes to communicate together. A same node can participate to several virtual rings at the same time. Each virtual ring is identified by a

unique identifier (named Virtual Ring Id). The Virtual Ring identifier is statically configured in all the nodes participating in the virtual ring. The way the virtual ring is initiated and managed will be described hereafter.

TCP/IP Protocol

- 5 In a preferred embodiment, the current invention is implemented on top of the TCP layer of the TCP/IP protocol, which is today the protocol the most largely used in the world. However, the invention only uses the transport function of TCP. It is also possible to implement the invention on top of any other protocol stack providing the transport function, such as IPX (Internetwork Packet Exchange). IP has been chosen
- 10 in the present description because this protocol is used in most of the networks. The transport function of TCP brings some reliability because this function handles transmission problems such as packet losses. The circulation of information along the virtual ring is based on the Internet Protocol (IP) and the Transmission Control Protocol (TCP). TCP has been chosen because it allows a sending of packets
- 15 without risk of loss. TCP also informs of the loss of the remote node by maintaining a connection. The use of TCP and IP allows to extend the virtual ring to any part of an IP network including the Internet itself. It is possible to imagine nodes in different parts of the world, communicating together with such a virtual ring.

The User Datagram Protocol (UDP) can also be used in the current invention for

- 20 instance to exchange Ring Insertion and Ring Removal messages between a specific node and the Virtual Ring Manager. Since these messages are exchanged only during the insertion or removal process, there is no need to use the TCP protocol and to establish a TCP session.

The present invention requires a new specific piece of code in each node part of the

- 25 ring network. This code uses a specific TCP port and a specific UDP port reserved for the invention. This code is used to establish, maintain and tear down the virtual ring topology

Token

In order to maintain the ring topology, some pieces of information need to be periodically exchanged between the different nodes. One of these pieces of information is called "token", referring to the "Token Ring" architecture developed by

5 IBM (IBM is a trademark of International Business Machines Corporation) these last decades. Figure 4 describes a token (401) circulating between node A and node B on a virtual ring (400).

The token is used as a periodic keepalive message to validate the ring topology. The token is periodically generated by the Virtual Ring Manager (402) and forwarded by

10 each node to its downstream neighbor node. The receipt by the Virtual Ring Manager of the token (from its upstream neighbor node), indicates that the ring topology is valid and the loop is not broken. If the ring is broken for some reason (loss of one node, loss of connectivity between 2 neighbor nodes), the loss of the token will indicate that there is a problem on the ring. Each node monitors the
15 reception of the token. If the token has not been received after a certain amount of time, each node will trigger the Ring Recovery process detailed here after. The token is forwarded from node to node, just like any other piece of information. This means that the Token uses the TCP sessions established between the nodes.

The Sequence Number field is used to identify the current copy of the token.

20 **Token structure**



The Token is described in Figure 5

- **IP header (500):** source IP address of the sending node and destination IP address of the next node in the virtual ring
- 25 • **TCP header (501):** source and destination ports = well known port reserved for the current invention
- **Virtual Ring Token (502):** This message contains 3 fields:

1. Message code (503), set to 0x0001. Allows to identify that the type of message is a Token.
2. Virtual Ring Identifier (504) on 2 bytes: identify the Virtual Ring. This allows a same node to participate to multiple Virtual Rings.
- 5 3. Sequence Number (505) on 4 bytes: it is set and incremented by the Virtual Ring Manager. This allows the Virtual Ring Manager to detect a possible duplication of the token.

Data Propagation along the virtual ring

When a node participating in the virtual ring receives a datagram from its upstream neighbor node, it processes this datagram, i.e stores the data part of the received message, and forwards it to its downstream neighbor node so that the datagram can circulate along the virtual ring. However, a node connected to the virtual ring must be able to recognize a datagram circulating along the virtual ring versus a normal IP datagram received from another node which does not participate in the virtual ring.

- 15 To do so, datagrams exchanged on the virtual ring have the following encapsulation:

IP Header (20 bytes)	TCP Header Source/Dest Port (20 bytes)	Virtual Ring Header Message Code 0x0000 Virtual Ring Identifier (2 bytes) Sender IP address (4 bytes)	Data
-------------------------	--	---	------

The encapsulation of the Data inside a TCP datagram has the following advantage : the datagram is transmitted along the Virtual Ring using the reliable TCP protocol.

The Virtual Ring Header comprises the following fields:

- 20 1. Message code : indicates that the received message is a datagram
2. Virtual Ring Identifier: indicates on which Virtual Ring the message must be forwarded. A node may belong to multiple Virtual Rings.
3. Sender IP address: This is the IP address of the node who has generated the data.

25 Transmission of a Datagram on the Virtual Ring

1/ When a node needs to send a datagram on the Virtual Ring, this node adds the Virtual Ring Header described above, and encapsulates the data inside a TCP datagram. This datagram is sent to the downstream neighbor on the Virtual Ring.

2/ Each node on the Virtual Ring checks the sender address to see which node has generated the datagram. Each node then reads the data, processes it, and forwards the datagram to its downstream neighbour.

3/ When the datagram is received back by the sender Node (the sender Node 5 checks the Sender IP address in the Virtual Ring Header), then the Sender Node removes the datagram from the Virtual Ring. This just means that the datagram is deleted and not forwarded to the downstream neighbor node again.

Virtual Ring Topology

The virtual ring is a list of nodes connected to form a ring. No node has the complete 10 view of the ring. This list of nodes participating in the ring is stored nowhere in the network. Each node comprises the following information (Node Ring Record) :

Virtual Ring Identifier (2 bytes) (configured)
Upstream Neighbor IP address (4 bytes)
Downstream Neighbor IP address (4 bytes)
Virtual Ring Manager IP address (configured) (4 bytes)
Backup Virtual Ring Manager IP address (configured) (4 bytes)

Virtual Ring Manager

One of the nodes participating in the virtual ring plays the role of "Virtual Ring 15 Manager". The Virtual Ring Manager is responsible for maintaining the topology of the virtual ring, more particularly the Virtual Ring Manager is responsible for the insertion and removal of the nodes.

It is important to note that the Virtual Ring Manager IP address is statically 20 configured in each node of the virtual ring. Since the Virtual Ring Manager constitutes a single point of failure, a Backup Virtual Ring Manager is generally used. The IP address of the Backup Virtual Ring Manager is also statically configured in each node. When a node wants to be inserted into the virtual ring and does not receive any response from the Virtual Ring Manager, this node will contact the Backup Virtual Ring Manager.

Insertion of a Node in the Virtual Ring

Figure 6 describes the insertion of a new node G (601) into a virtual ring (600) comprising nodes A, B, C, D, E, F. When a new node G (601) wants to join the virtual ring (600), the following scenario occurs:

5 **Note:** in a preferred embodiment, all the insertion messages use the UDP protocol and the reserved UDP port defined in the current invention.

- The Node (601) to insert in the virtual ring (Node G), sends a "Virtual Ring Insertion Request" message (603) to the Virtual Ring Manager (602) using the configured IP address of the Virtual Ring Manager. The Node (601) to insert 10 starts an "Insertion Request" timer and waits for a "Virtual Ring Insertion Confirmation" message (604).
- The Virtual Ring Manager (602) receives the "Virtual Ring Insertion Request" message and notes the source IP address of the message, which is the IP address of Node G (601).
- The Virtual Ring Manager (602) sends a "Virtual Ring Change Neighbor" message (605) to its downstream neighbor Node F (606). The Virtual Ring Manager finds the IP address of Node F in its Node Ring Record. The "Virtual Ring Change Neighbor" message comprises the IP address of Node G (601) as Upstream Neighbor IP address. The Downstream Neighbor IP address in the 15 message is set to 0.0.0.0 because this address does not need to be changed.
- Node F (606) receives the "Virtual Ring Change Neighbor" message (605). Node F tears the TCP session down with its upstream neighbor Node (the Virtual Ring Manager), by issuing a "TCP Reset" message. Node F (606) stores the IP address of Node G (601) received in the "Virtual Ring Change Neighbor" 20 Message (605), in its Node Ring Record (Upstream Neighbor IP address).
- Node F (606) establishes a TCP session with its new upstream neighbor Node, (Node G (601)) and sends a "Virtual Ring Neighbor Changed" message (607) to the Virtual Ring Manager (602) to indicate that Node F has changed its upstream neighbor Node.
- The Virtual Ring Manager (602) receives the "TCP Reset" message from Node F (606) and tears the TCP session down. The Virtual Ring Manager establishes a 25

new TCP session with Node G (601) and stores the IP address of Node G (601) in its Node Ring Record: Downstream Neighbor IP address.

- The Virtual Ring Manager (602) sends a "Virtual Ring Insertion Confirmation" message (604) to Node G (601). This message comprises the IP address of Node F (606).
- Node G (601) updates its Node Ring Record with :
 - an Upstream Neighbor IP address equal to the Virtual Ring Manager IP address, and
 - an Downstream Neighbor IP address equal to the IP address of Node F.

10 Node G (601) stops the "Insertion Request" timer.

- If the "Insertion Request" timer expires, this means that Node G (601) has not received the "Virtual Ring Insertion Confirmation" message (604) from the Virtual Ring Manager (602). In that case, Node G (601) contacts the Backup Virtual Ring Manager (608). This process is described below in the section related to the

15 Backup Virtual Ring Manager.

The result of the insertion of node G is described in Figure 7. Node G (701) is now inserted on the virtual ring (700), between the Virtual Ring Manager (702) and Node F (703).

Solicited Removal of a Node from the Virtual Ring

20 The solicited node removal scenario described in the present section correspond to the case where a node wants to be removed from the Virtual Ring because it does not want to participate any more in the group.

Another node removal scenario corresponds to the case where a node has a failure and the virtual ring is broken. This unsolicited removal scenario will be described in
25 another section.

Figure 8 describes the Node Solicited removal process. When Node C (801) wants to be removed from the virtual ring, the following scenario occurs:

- Node C (801) sends a "Virtual Ring Removal Request" message (803) to the Virtual Ring Manager (802). This message comprises
 - the IP address of Node C (801),

- the IP address of its upstream neighbor node, Node B (804), and
- the IP address of its downstream neighbor node, Node D (805).

Node C (801) starts a "Ring Removal" Timer and waits for a "Virtual Ring Removal Confirmation" message (806)

- 5 • The Virtual Ring Manager (802) receives the "Virtual Ring Removal Request" message (803) from Node C (801) and starts the removal process. It notes the IP addresses of the upstream neighbor node and downstream neighbor node of Node C.
- 10 • The Virtual Ring Manager (802) sends a "Virtual Ring Change Neighbor" message (807) to Node B (804), upstream neighbor node of the node to remove, Node C (801). This message comprises:
 - the unchanged Upstream Node IP address: 0.0.0.0;
 - the Downstream Node IP address equal to the IP address of Node D (805), which is the downstream neighbor node of Node C (801)
- 15 • The Virtual Ring Manager (802) sends a "Virtual Ring Change Neighbor" message (808) to Node D (805), downstream neighbor node of the node to remove, Node C (801). This message comprises:
 - the Upstream Node IP address equal to the IP address of Node B (804), which is the upstream neighbor node of Node C (801);
 - the unchanged Downstream Node IP address: 0.0.0.0.

The Virtual Ring Manager (802) starts a "Change Neighbor" Timer of, for instance, 30s waiting for the confirmations.

- Node B (804) receives the "Virtual Ring Change Neighbor" message (807) from the Virtual Ring Manager (802) and modifies its Node Ring Record.
- 25 The Upstream Node IP address in the message is 0.0.0.0. This means that the address does not need to be changed. Node B (804) keeps Node A (809) as its Upstream Neighbor node.
On the other hand, Node B (804) modifies its Downstream Neighbor IP address and uses the address received in the message. Node B (804) tears the TCP connection down with Node C (801) which was its previous downstream neighbor node, and establishes a new TCP connection with its new downstream neighbor node, Node D (805).
- 30 Node B (804) sends a "Virtual Ring Neighbor Changed" message (810) to the Virtual Ring Manager (802).

- Node D (805) does the same as Node B (804). It updates its Node Ring Record, tears the TCP session it had with its Upstream Neighbor, Node C (801) down, and establishes a TCP session with its new Upstream Neighbor, Node B (804).
5 Node D (805) sends a "Virtual Ring Neighbor Changed" message (811) to the Virtual Ring Manager (802)
- When the Virtual Ring Manager receives the "Virtual Ring Neighbor Changed" messages from both nodes B and C, it stops the "Change Neighbor" timer and sends a "Virtual Ring Removal Confirmation" message (806) to node C (801) to indicate that the Removal process has been successfull.
- 10 • Node C (801) stops the "Ring Removal" Timer. If the "Ring Removal" Timer expires, it means that the Virtual Ring Manager (802) has not achieved the Removal process. In this case, Node C (801) must contact the Backup Virtual Ring Manager (804).

Loss of a Node

- 15 The loss of a node in the virtual ring network is detected by its neighbor nodes with the loss the TCP connections. When a node is removed from the virtual ring without informing the Virtual Ring Manager by means of a "Virtual Ring Removal Request" message (which should be the case when a node failure occurs), the 2 neighbor nodes (upstream and downstream) lose their TCP connection with this node a given period of time (after a TCP timeout). As described in Figure 9; the following scenario occurs:

- Node C (901) in the virtual ring network (900), fails or is powered off.
- Node B (903), the upstream neighbor node of Node C (901), loses its TCP connection with Node C. Node B attempts to re-establish its TCP connection without success.
- 25 • Node D (904), downstream neighbor node of Node C (901), loses its TCP connection with Node C. Node D attempts to re-establish its TCP connection without success.
- Node B (903) sends a "Virtual Ring Neighbor Loss Indication" (907) message to the Virtual Ring Manager (902). This message comprises:
 - 30 the Node B IP address, and
 - the Node B Downstream Neighbor IP address, i.e. address of Node C (901).

The Upstream Neighbor IP address is set to 0.0.0.0 because no problem has been found with the upstream neighbor node of Node B.

- Node D (904) sends a "Virtual Ring Neighbor Loss Indication" (905) message to the Virtual Ring Manager (902). This message comprises:

5 • the Node D IP address, and

 • the Node D Upstream Neighbor IP address, i.e. address of Node C (901).

The Downstream Neighbor IP address is set to 0.0.0.0 because no problem has been found with downstream neighbor node of Node D.

- The Virtual Ring Manager (902) receives both messages from Node C upstream neighbor node and Node C downstream neighbor node. The virtual ring needs to be reconfigured.

- The Virtual Ring Manager (903) sends a "Virtual Ring Change Neighbor" message (906) to Node D (904). This message comprises:

 • an Upstream Neighbor IP address equal to the Node B IP address

15 • an unchanged Downstream Neighbor IP address equal to 0.0.0.0 (the IP address does not need to be changed)

- Node D (904) updates the Upstream Neighbor IP address in its Node Ring Record and establishes a TCP connection with its new upstream neighbor node (Node B (903)).

20 • Node D (904) sends a "Virtual Ring Node Changed" (909) message back to the Virtual Ring Manager (902) to confirm the change.

- The Virtual Ring Manager (903) sends a "Virtual Ring Change Neighbor" message (908) to Node B (903) comprising:

 • an unchanged Upstream Neighbor IP address equal to 0.0.0.0 (the IP address does not need to be changed)

25 • a Downstream Neighbor IP address equal to the Node D IP address.

- Node B (903) updates the Downstream Neighbor IP address in its Node Ring Record.

30 • Node B (903) sends a "Virtual Ring Node Changed" message (910) back to the Virtual Ring Manager (902) to confirm the change.

Figure 10 shows the result of the virtual ring (1000) reconfiguration after the loss of Node C (1001).

Backup Virtual Ring Manager

The Backup Virtual Ring Manager executes the same processes as the Virtual Ring Manager. The Backup Virtual Ring Manager receives Insertion, Removal and Recovery messages from the nodes in absence of response from the Virtual Ring

5 Manager, and processes these messages like the Virtual Ring Manager.

Token Loss Recovery

All the nodes including the Virtual Ring Manager, use a timer to detect the loss of the token. When the token is lost, the ring needs to be rebuilt. The value of this timer must be larger than the TCP session timer to allow the process described in section

10 entitled "Loss of a node" to take place before the reconfiguration of the ring. When a node detects the loss of the token, it sends a "Virtual Ring Removal Request" message to the Virtual Ring Manager and waits for the confirmation as described in Figure 8 (refer to section entitled "Solicited Node Removal"). After a given period of time, the node will send a "Virtual Ring Insertion Request" message to the Virtual

15 Ring Manager to participate again in the ring as described in Figure 6 (section entitled "Insertion of a Node").

INSERTION AND REMOVAL MESSAGES

These messages are exchanged using the User Datagram Protocol (UDP). The value of the Virtual Ring Identifier field is used to identify the current virtual ring. The

20 Virtual Ring Identifier is statically configured in each participating node.

General Format



Virtual Ring Insertion Request



25 Virtual Ring Insertion Confirmation

Message Code 0x0003	Virtual Ring Identifier (2 bytes)	Upstream Neighbor IP address (4 bytes)	Downstream Neighbor IP address (4 bytes)
------------------------	--------------------------------------	---	---

Virtual Ring Change Neighbor

Message Code 0x0004	Virtual Ring Identifier (2 bytes)	Upstream Neighbor IP address (4 bytes)	Downstream Neighbor IP address (4 bytes)
------------------------	--------------------------------------	---	--

Virtual Ring Neighbor Changed

Message Code 0x0005	Virtual Ring Identifier (2 bytes)	Upstream Neighbor IP address (4 bytes)	Downstream Neighbor IP address (4 bytes)
------------------------	--------------------------------------	---	--

5

Virtual Ring Removal Request

Message Code 0x0006	Virtual Ring Identifier (2 bytes)	Removing Node IP address (4 bytes)	Upstream Neighbor IP address (4 bytes)	Downstream Neighbor IP address (4 bytes)
------------------------	---	---	---	---

Virtual Ring Removal Confirmation

Message Code 0x0007	Virtual Ring Identifier (2 bytes)
------------------------	---

10 Virtual Ring Neighbor Loss Indication

Message Code 0x0008	Virtual Ring Identifier (2 bytes)	Upstream Neighbor IP address (4 bytes)	Downstream Neighbor IP address (4 bytes)	Node IP address (4 bytes)
------------------------	---	---	---	---------------------------------

PROCESSES ACCORDING TO THE PRESENT INVENTION

Token processing in a Node

Figure 11 describes the algorithm executed by a node when this node receives the

15 Token.

- (1100) The Node has just been inserted into the virtual ring.
- (1101) The Node starts the Wait Token Timer (30 seconds) and waits for the receipt of the Token
- (1102) The Node checks whether the Token has been received or not.

- (1103) If no Token has been received, the Node checks whether the Token Timer has expired or not. If the Token Timer has not expired, the Node continues to wait for the Token.
- (1104) If the Token has been received, the Node checks the Token Sequence number to verify that it has been incremented since the last reception. If the Token is received for the first time (just after the node insertion), this test is not executed.
- 5 • (1105) If the Token Sequence number in the received Token is correct, the Node forwards the Token to its downstream neighbor node and waits for the Token again.
- 10 • (1106) If no Token has been received and if the Token Timer has expired, or if the received Token do not have the expected Token Sequence number (this means that a Token has been lost), then the Ring Recovery Procedure is executed.

15 Token processing in the Virtual Ring Manager

Figure 12 describes the algorithm executed in the Virtual Ring Manager at receipt of the Token.

- (1200) The Virtual Ring Manager has just been inserted. It sets the Token Sequence number to 1, starts a Wait Token Timer of 30 seconds and a Token Timer of 1 second. The Token Timer is used to generate a new Token every second. The Wait Token Timer is used to trigger the ring recovery.
- 20 • (1201) The Virtual Ring Manager forwards the Token to its downstream neighbor node and waits for the return of the Token.
- (1202) The Virtual Ring Manager checks whether the Token has been received or not.
- 25 • (1203) If the Token has not been received, the Virtual Ring Manager checks whether the Token Timer has expired or not.
- (1204) If the Token Timer has not expired, the Virtual Ring Manager checks whether the Wait Token Timer has expired or not. If not, the Virtual Ring Manager waits for the Token again.
- 30

- (1205) If no Token has been received and if the Wait Token Timer has expired, this means that the Token has been lost. Then the Virtual Ring Manager executes the ring recovery procedure.
- (1206) If the token is received, the Virtual Ring Manager checks the sequence number in the Token.
- 5 • (1207) The Virtual Ring Manager restarts the Wait Token Timer because the token has been received and wait for the Token timer to expire.
- (1208) When the Token Timer expires, the Virtual Ring Manager generates a new token, and increments the sequence number
- 10 • (1209) The Virtual Ring Manager forwards the Token to its downstream neighbor node and wait for the return of the Token.

Node Insertion – Algorithm in the Node

Figure 13 describes the algorithm executed in a node in view of inserting this node into the virtual ring.

- 15 • (1300) A new Node wants to be inserted into the Virtual Ring. This Node sends a "Virtual Ring Insertion Request" message to the Virtual Ring Manager.
- (1301) the Node starts the Insertion Timer and waits for an "Virtual Ring Insertion Confirmation" message from the Virtual Ring Manager.
- (1302) the Node checks whether an "Virtual Ring Insertion Confirmation"
- 20 message has been received or not.
- (1303) If no "Virtual Ring Insertion Confirmation" message has been received, the Node checks if the Insertion Timer has expired. If not, the Node continues to wait for the receipt of an "Virtual Ring Insertion Confirmation" message.
- (1304) If an "Virtual Ring Insertion Confirmation" message has been received, the Node stops the Insertion Timer, updates its neighbor addresses and establishes
- 25 TCP sessions with its neighbor nodes.
- (1305) the new Node is now inserted into the virtual ring.
- (1306) If no "Virtual Ring Insertion confirmation" message has been received and if the Insertion Timer has expired, the Node sends an "Virtual Ring Insertion
- 30 Request" message to the Backup Manager.

- (1307) the Node starts the Insertion Timer again.
- (1308) The Node checks whether an "Virtual Ring Insertion Confirmation" message has been received or not.
- (1309) If no "Virtual Ring Insertion Confirmation" message has been received, the Node checks whether the Insertion Timer has expired or not. If not, the Node continues to wait for the receipt of an "Insertion Confirmation" message.
- 5 • (1310) If an "Virtual Ring Insertion Confirmation" message has been received, the Node stops the Insertion Timer, updates its neighbor addresses and establishes TCP sessions with its neighbor nodes.
- 10 • (1311) the new Node is now inserted into the virtual ring.
- (1312) if no "Virtual Ring Insertion Confirmation" message has been received from the Backup Manager, this means that both the Virtual Ring Manager and the Backup Manager are unavailable. In this case, the process for inserting the Node in the ring has failed.

15 Node Removal – Algorithm in the Node

Figure 14 describes the algorithm executed in a node in view of removing this node from the virtual ring.

- (1401) A Node wants to be removed from the virtual ring. This Node sends a "Virtual Ring Removal Request" message to the Virtual Ring Manager
- 20 • (1402) The Node to remove starts the Ring Removal Timer.
- (1403) The Node to remove waits for a "Virtual Ring Removal Confirmation" message from the Ring Manager.
- (1404) When the "Virtual Ring Removal Confirmation" message is received, the Node terminates its shutdown.
- 25 • (1405) If no "Virtual Ring Removal Confirmation" message is received, the Node to remove checks the Ring Removal Timer.
- (1406) If the Ring Removal Timer has expired, the Node to remove sends a "Virtual Ring Removal Request" message to the Backup Ring Manager
- (1407) The Node to remove starts the Ring Removal Timer.
- 30 • (1408) The Node to remove waits for the "Virtual Ring Removal Confirmation" message from the Ring Manager.

- (1409) If no "Virtual Ring Removal Confirmation" message is received, the Node to remove checks the Ring Removal Timer. If this timer expires, then the Node to remove terminates its shutdown without waiting for any response from the Ring Manager.

5 Node Insertion/Removal – Algorithm in the Adjacent Node

Figure 15 illustrates the algorithm executed in a node when a neighbor node has been inserted or removed.

- (1501) The Adjacent Node checks if a "Virtual Ring Change Neighbor" message has been received from the Virtual Ring Manager.
- 10 • (1502) If a "Virtual Ring Change Neighbor" message has been received, the Adjacent Node updates its neighbor table using the upstream and downstream addresses received in the message.
- (1503) The Adjacent Node sends back a "Virtual Ring Neighbor Changed" message to the Virtual Ring Manager.

15 Node Insertion/Removal – Algorithm in the Virtual Ring Manager

Figure 16 illustrates the algorithm executed in the Virtual Ring Manager when a node is inserted or removed from the virtual ring.

- (1601) The Virtual Ring Manager checks if a "Virtual Ring Insertion Request" message has been received.
- 20 • (1602) If a "Virtual Ring Insertion Request" message has been received, the Virtual Ring Manager sends a "Virtual Ring Change Neighbor" message to its own downstream neighbor, and starts a Change Neighbor Timer.
- (1603) the Virtual Ring Manager waits for a "Virtual Ring Neighbor Changed" message. If no "Virtual Ring Neighbor Changed" message is received and the Change timer expires, then the procedure fails.
- (1604) the Virtual Ring Manager updates its downstream address with the address of the new Node.
- (1605) The Virtual Ring Manager sends an "Virtual Ring Insertion Confirmation" message to the Node to insert.

- (1606) The Virtual Ring Manager checks if a "Virtual Ring Removal Request" message has been received.
- (1607) The Virtual Ring Manager sends a "Virtual Ring Change Neighbor" message to the downstream neighbor of the Node to remove.

5 • (1608) The Virtual Ring Manager sends a "Virtual Ring Change Neighbor" message to the upstream neighbor of the Node to remove.

- (1609) The Virtual Ring Manager starts the Change Neighbor Timer.
- (1610) When the "Virtual Ring Neighbor Changed" messages have been received from the upstream and the downstream neighbor nodes, the Virtual Ring

10 Manager sends a "Virtual Ring Removal Confirmation" message to the Node to remove.

- (1611) The Virtual Ring Manager checks if a "Virtual Ring Neighbor Loss Indication" message has been received.

Flow of Node Insertion

15 Figure 17 illustrates the Node Insertion process.

- (1701) Node G is the Node to insert. Its ring table contains the Virtual Ring and Backup Ring Manager addresses.
- (1702) Node E is the Virtual Ring Manager.
- (1703) Node F is the downstream neighbor node of the Virtual Ring Manager.

20 • (1704) Inserting Node G sends a "Virtual Ring Insertion Request" message to the Virtual Ring Manager.

- (1705) The Virtual Ring Manager sends a "Virtual Ring Change Neighbor" message to its downstream neighbor node (node F) in order to insert the new Node G just before Node F.

25 • (1706) Node F updates its Virtual Ring table.

- (1707) Node F replies with a "Virtual Ring Neighbor Changed" message to the Virtual Ring Manager.
- (1708) The Virtual Ring Manager updates its Ring Table to insert the new Node G just after itself.

30 • (1709) The Virtual Ring Manager sends a "Virtual Ring Insertion Confirmation" message to the new inserting Node G.

- (1710) Node G updates its Virtual Ring table with the addresses of 2 adjacent nodes.

Flow of Solicited Node Removal

Figure 18 illustrates the Solicited Node Removal process.

- 5 • (1801) Node B is the upstream neighbor node of Node C, the Node to remove.
- (1802) Node C is the Node to remove.
- (1803) Node D is the downstream neighbor node of Node C.
- (1804) Node E is the Virtual Ring Manager.
- (1805) Removing Node C sends a "Virtual Ring Removal Request" message to the Virtual Ring Manager.
- 10 • (1806) The Virtual Ring Manager sends a "Virtual Ring Change Neighbor" message to Node D (downstream neighbor node).
- (1807) The Virtual Ring Manager sends a "Virtual Ring Change Neighbor" to Node B (upstream neighbor node) and starts the Change Neighbor Timer.
- 15 • (1808) Node B updates its Ring Table.
- (1809) Node B sends a "Virtual Ring Neighbor Changed" message to the Virtual Ring Manager.
- (1810) Node D updates its Ring Table.
- (1811) Node D sends a "Virtual Ring Neighbor Changed" message to the Virtual Ring Manager.
- 20 • (1812) Virtual Ring Manager stops the Change Neighbor Timer and sends a "Virtual Ring Removal Confirmation" message to the Node C to remove.

Flow of Unsolicited Node Removal (node loss)

Figure 19 illustrates the unsolicited Node Removal process.

- 25 • (1901) Node B is the upstream neighbor node of Node C, the Node to remove.
- (1902) Node C is the Node to remove.
- (1903) Node D is the downstream neighbor node of Node C.
- (1904) Node E is the Virtual Ring Manager.

- (1905) The upstream neighbor node B detects the loss of the TCP connection with node C and sends a "Virtual Ring Loss Indication" message to the Virtual Ring Manager.
- (1906) The downstream neighbor node D detects the loss of the TCP connection with node C and sends a "Virtual Ring Loss Indication" message to the Virtual Ring Manager.
- 5 • (1907) The Virtual Ring Manager sends a "Virtual Ring Change Neighbor" message to Node D (downstream neighbor node).
- (1908) The Virtual Ring Manager sends a "Virtual Ring Change Neighbor" message to Node B (upstream neighbor node) and starts the Change Neighbor Timer.
- 10 • (1909) Node B updates its Ring Table.
- (1910) Node B sends a "Virtual Ring Neighbor Changed" message to the Virtual Ring Manager.
- 15 • (1911) Node D updates its Ring Table.
- (1912) Node D sends a "Virtual Ring Neighbor Changed" message to the Virtual Ring Manager.

ADVANTAGES OF THE PRESENT INVENTION

1) Reduction of the number of TCP sessions in the network.

20 With the present invention, only N TCP sessions are required to interconnect N nodes versus $Nx(N-1)/2$ sessions in a full meshed network or $2xN-3$ sessions in a dual star configuration. A session is a virtual connection between two nodes, enabling the exchange of data between these nodes and taking care of transmission problems like flow control and retransmission. A TCP session is an example of

25 session between two nodes supporting the TCP/IP protocol. The current invention is implemented on top of the TCP (Transmission Control Protocol) layer of the TCP/IP protocol stack, and can be used by any node supported the TCP/IP protocol, which is the protocol the most widely used in the world. As illustrated in Figure 1, in a full meshed network of N nodes, each node has to establish a TCP session with each of

30 the N-1 other nodes. This means the establishment of $Nx(N-1)/2$ TCP sessions.

According to the present invention, each node has to establish a TCP session with only 2 other nodes: the upstream neighbor node and the downstream neighbor node.

This means a total of N sessions in a virtual ring of N nodes. The saving resulting from the present invention can be calculated as follows :

$Nx(N-1)/2$ sessions in a full meshed network versus N sessions in the virtual ring configuration. The difference is equal to $Nx(N-1)/2-N = N^2/2-N/2-N= N^2/2-3N/2=$

5 $Nx(N-3)/2$. Therefore, the current invention allows to save $Nx(N-3)/2$ sessions in the network.

Reducing the number of sessions between the nodes brings several other advantages as it will be explained in the following points.

2) Reduction of the bandwidth utilisation in the network

10 Because each node receives one and only one copy of a same message, the present invention avoids multiple and unnecessary copies. In a full meshed topology as described in Figure 1, each node communicates with all the other nodes. If each node needs to send the same piece of information to the other nodes, each node will forward this piece of information to all its neighbor nodes, and this will duplicate the
15 number of messages exchanged in the network. This is typically the case when the nodes are routers exchanging routing information using a routing protocol like RIP (Routing Information Protocol). Periodically, each router communicates, or floods, its routing table to the other routers in the network. Another example is when a distributed database needs to be synchronised, and when the servers participating in
20 the distributed database need to exchange a same record. Usually, the broadcast of information is managed by the application layer, which must take care of the way the information is distributed between the nodes.

The present invention enable the exchange of a same piece of information between all the nodes so that each node receives one and only one copy of the information.

25 Because nodes are virtually connected to a virtual ring, and because the information circulates along that ring and is seen by each node connected to the ring, the network is not flooded by multiple copies of messages exchanged between nodes.

3) Reduction of the CPU utilization in each node.

Establishing and Maintaining a TCP session requires computer resources to manage
30 the flow control, the retransmission, and to generate acknowledgements and keepalives. The present invention reduces the number of TCP sessions required for

nodes to communicate, and therefore reduces the utilization of data processing resources in nodes.

4) Reduction of the memory consumption inside each node

In each node, the maintenance of TCP sessions requires to keep the context of 5 these session, with information such as the sequence number of the last segment sent, or the sequence number of the next acknowledgement to send. The storage of this information consumes memory. Reducing the number of TCP sessions has for result to reduce the memory consumption in the nodes.

While the invention has been particularly shown and described with reference to a 10 preferred embodiment, it will be understood that various changes in form and detail may be made therein without departing from the spirit, and scope of the invention.